

# Ethernet over Kernel Sockets

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-02/msg15301.html>

---

- *From:* Lukas Razik <[linux@xxxxxxxxxx](mailto:linux@xxxxxxxxxx)>
  - *Date:* Fri, 29 Feb 2008 21:54:21 +0100
- 

Hello all!

As you know, some network cards doesn't have a 'eth' interface under linux. Because of that I'm developing a net\_device based driver which doesn't transmit and receive directly through a real network card but through UDP kernel sockets.

That means:

If my net\_device->hard\_start\_xmit function gets an packet to transmit (in interrupt context), there will be a work struct queued into a workqueue and the packet will be processed by the workqueue later on (in process context).

On the receiver side I have a kernel thread which blocks on sock\_recvmsg() and if there comes a UDP message, then it will be processed and a sk\_buff will be passed to the kernel.

The current state is that all works stable but I've bad transmission rates and bad ping times.

For example I've to systems and each has one Gigabit-Ethernet card:

System 1: ifconfig eth0 192.168.0.1

System 2: ifconfig eth0 192.168.0.2

If I load my driver then I get an additional eth interface (for example):

System 1: ifconfig eth1 192.168.1.1

System 2: ifconfig eth1 192.168.1.2

So the eth1 interfaces are based on sockets which use the eth0 (Gigabit) interfaces for the communication.

Now, if I measure the transmission between 192.168.0.1 and 192.168.0.2 then I get transmission rates that are normal for Gigabit-Ethernet (~25µsec (PingPong) and ~900Mbit/s).

If I measure the transmission through the eth1 interfaces between 192.168.1.1 and 192.168.1.2 then I only get ~1msec (PingPong) and only about ~400Mbit/s.

The interesting thing is, that I always get PingPongs of 1ms (= one jiffie). Although if I measure with 100MBit- or other non-Ethernet cards.

Maybe it's because I run the socket\_recvmsg() function in a kernel thread which blocks on the function and must wait to be executed by the scheduler.

Now I don't know how I could solve this problem:

How can I force the kernel to process a received message immediately when it comes in through an UDP kernel socket?

I hope that someone of you can help me....

The source code you can find here:

<http://net.razik.de/ethos.tar.gz>

## Ethernet over Kernel Sockets

Regards and Many Thanks for any help!

Lukas

--

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>